

What is claimed is:

What is claimed is:

- a plurality of input ports for partitioning portions of received data frames to provide data cells;

a plurality of switching sections, each switching section being coupled to each input port for receiving data cells at cell transfer intervals on a data link coupled between the input port and switching section, the switching section being coupled to transmit data cells to any one of a plurality of output ports,

wherein each input port includes logic for scheduling the transmission of each data cell of each data frame received at the input port, the data frame having a destination associated with an output port, during a cell transfer interval for each data link coupled between the input port and each switching section based upon an ability of the switching section to receive data cells of data frames with a destination associated with the output port.

2. The switching fabric of claim 1, wherein each of the input ports maintains a plurality of data frame queues of received data frames, each of the data frame queues corresponding with one of the output ports and having logic for enqueueing data frames having a destination associated with the output port.

3. The switching fabric of claim 2, wherein each of the data frames includes a data payload and the input ports provides for each data frame one or more associated data cells including a

10

15

portion of the data payload, the one or more data cells associated with the data frame collectively having the data payload of the data frame, wherein each input port schedules a transmission of each data cell to one of the switching sections on the data link coupled between the input port and the switching section.

5

4. The switching fabric of claim 3, wherein each input port schedules a transmission of each data cell to one of the switching sections on the data link coupled between the input port and the switching section.

SUB 417

10

5. The switching fabric of claim 3, wherein for each data link coupled between each input port and each switching section, the input port attempts to schedule a data cell of a partially transmitted data frame, the partially transmitted data frame having at least one associated data cell previously scheduled for transmission to a switching section, from any of the data frame queues subject to the ability of the switching section to receive data cells of data frames having a destination associated with the output port associated with the destination of the partially transmitted data frame prior to scheduling a transmission of a data cell of a data frame for which no data cells have been previously scheduled for transmission to a switching section.

15

6. The switching fabric of claim 1, wherein each of the switching sections maintains a plurality of data cell queues of data cells received on the data links coupling the switching section to the input ports, each of the data cell queues corresponding with an output port, each of the data cells in the data cell queue being of a partition of a portion of a data frame having a destination associated with the output port.

20

7. The switching fabric of claim 6, wherein each of the data cell queues of a switching section is capable of enqueueing a finite number of data cells at any one time, and wherein the ability of a switching section to receive data cells of data frames with a destination associated with an output port is based upon a quantity locations in the data cell queue which are capable of receiving a data cell from an input port.

8. The switching fabric of claim 1, the switching fabric further including a plurality of output ports, each output port having logic for reassembling data frames having a destination associated with the output port from data cells received from the switching sections coupled to the output port.

9. The switching fabric of claim 8, wherein each output is coupled to one or more media access control (MAC) devices through a common transmission medium, and wherein for each MAC device coupled to the output port, the output port maintains an associated MAC queue of reassembled data frames for transmission to the MAC device, the destination of each reassembled data frame in the MAC queue being associated with the MAC device.

10. The switching fabric of claim 9, wherein each of the output ports transmits a signal to each of the switching sections indicating an ability to receive data cells from data links coupling the output port to the switching section.

11. The switching fabric of claim 1, wherein the switching fabric includes a plurality of output ports and for each of the output ports, each of the switching sections transmits a signal to each of the input ports indicating the ability of the switching section to receive data cells of data frames having a destination associated with the output port.

5

12. A method of transmitting digital data from a plurality of sources to a plurality of destinations, the method comprising:

receiving data frames at each of a plurality of input ports;

partitioning portions of received data frames to provide data cells;

receiving data cells at each of a plurality of switching sections at cell transfer intervals on a data link coupled between the switching section and an input port; and

transmitting data cells from each switching section to any one of a plurality of output ports; and

scheduling the transmission of each data cell of each data frame received at each the input port, the data frame having a destination associated with an output port, during a cell transfer interval for each data link coupled between the input port and each switching section based upon an ability of the switching section to receive data cells of data frames with a destination associated with the output port.

20

13. The method of claim 12, the method further comprising maintaining a plurality of data frame queues of received data frames at each of the input ports, each of the data frame queues corresponding with one of the output ports and enqueueing data frames having a destination associated with the output port.

14. The method of claim 13, wherein each of the data frames includes a data payload, the method further comprising:

providing for each data frame in a data frame queue at an input port one or more associated data cells including a portion of the data payload of the data frame, the one or more data cells associated with the data frame collectively having the data payload of the data frame; and

scheduling a transmission of each data cell to one of the switching sections on the data link coupled between the input port and the switching section.

15. The method of claim 14, the method further comprising scheduling a transmission of each data cell to one of the switching sections on the data link coupled between the input port and the switching section prior to scheduling a transmission of a data cell of a subsequent data frame in the data frame queue to any of the switching sections.

16. The method of claim 15, the method further comprising, for each data link coupled between each input port and each switching section, attempting to schedule a transmission of a data cell of a partially transmitted data frame, the partially transmitted data frame having at least one associated data cell previously scheduled for transmission to a switching section, from any of the data frame queues subject to the ability of the switching section to receive data cells of data frames having a destination associated with the output port associated with the destination of the partially transmitted data frame prior to scheduling a transmission of a data cell of a data

frame for which no data cells have been previously scheduled for transmission to a switching section.

17. The method of claim 12, the method further comprising, at each switching section,
5 maintaining a plurality of data cell queues of data cells received on the data links coupling the switching section to the input ports, each of the data cell queues corresponding with an output port, each of the data cells in the data cell queue being of a partition of a portion of a data frame having a destination associated with the output port.

10 18. The method of claim 17, wherein each of the data cell queues of a switching section is capable of enqueueing a finite number of data cells at any one time, the method further including determining the ability of the switching section to receive data cells of data frames with a destination associated with an output port is based upon a quantity locations in the data cell queue which are capable of receiving a data cell from an input port.

15 19. The method of claim 12, the method further comprising:

receiving data cells at each of a plurality of output ports from the switching sections coupled to the output ports; and

20 at each output port, reassembling data frames having a destination associated with the output port from data cells received from the switching sections coupled to the output port.

20. The method of claim 19, the method further comprising, at each output port, maintaining a media access control (MAC) queue of reassembled data frames to be transmitted to one or

SUBP 17

more MAC devices through a common transmission medium, the destination of each reassembled data frame in the MAC queue being associated with the MAC device.

21. The method of claim 17, the method further comprising transmits a signal from each of the output ports to each of the switching sections indicating an ability to receive data cells from data links coupling the output port to the switching section.

22. The method of claim 12, wherein the switching fabric includes a plurality of output ports and for each of the output ports, the method further comprising transmitting a signal from each switching section to each of the input ports indicating the ability of the switching section to receive data cells of data frames having a destination associated with the output port.

23. In a data communication network including a plurality of host computers for transmitting data packets to a plurality of destinations, each destination being associated with a media access control (MAC) device having a MAC address, the improvement including:

a plurality of output ports, each of the output ports being coupled to at least an associated one of the MAC devices for transmitting MAC data frames to the at least one MAC device according the MAC address associated therewith;

a look-up engine for receiving the data packets from the host computers and forming intermediate data frames based upon the data packets, the intermediate data frames having information identifying an output port associated with one of the destinations the network device in a header and a data payload;

SUB 94

10
15
20

a plurality of input ports for receiving the intermediate data frames from the lookup engine, each of the plurality of input ports partitioning the data payload of at least some of the intermediate frames received at the input port to provide a plurality of data cells;

a plurality of switching sections, each switching section being coupled to each input port
 5 for receiving data cells at cell transfer intervals on a data link coupled between the input port and switching section, the switching section being coupled to transmit data cells to any one of the plurality of output ports,

wherein each input port includes logic for scheduling the transmission of each data cell of each intermediate data frame received at the input port during a cell transfer interval for each
 10 data link coupled between the input port and each switching section based upon an ability of the switching section to receive data cells of data frames associated with the output port.

24. The data communication network of claim 23, wherein each of the input ports maintains a plurality of data frame queues of received data frames, each of the data frame queues
 15 corresponding with one of the output ports and enqueueing data frames having a destination associated with the output port.

25. The data communication network of claim 24, wherein each of the data frames includes a data payload and the input ports provides for each data frame one or more associated data cells
 20 including a portion of the data payload, the one or more data cells associated with the data frame collectively having the data payload of the data frame, wherein each input port schedules a transmission of each data cell to one of the switching sections on the data link coupled between the input port and the switching section.

26. The data communication network of claim 25, wherein each input port schedules a transmission of each data cell to one of the switching sections on the data link coupled between the input port and the switching section prior to scheduling a transmission of a data cell of a subsequent data frame in the data frame queue to any of the switching sections.

27. The data communication network of claim 25, wherein for each data link coupled between each input port and each switching section, the input port attempts to schedule a data cell of a partially transmitted data frame, the partially transmitted data frame having at least one associated data cell previously scheduled for transmission to a switching section, from any of the data frame queues subject to the ability of the switching section to receive data cells of data frames having a destination associated with the output port associated with the destination of the partially transmitted data frame prior to scheduling a transmission of a data cell of a data frame for which no data cells have been previously scheduled for transmission to a switching section.

28. The data communication network of claim 23, wherein each of the switching sections maintains a plurality of data cell queues of data cells received on the data links coupling the switching section to the input ports, each of the data cell queues corresponding with an output port, each of the data cells in the data cell queue being of a partition of a portion of a data frame having a destination associated with the output port.

29. The data communication network of claim 28, wherein each of the data cell queues of a switching section is capable of enqueueing a finite number of data cells at any one

time, and wherein the ability of a switching section to receive data cells of data frames with a destination associated with an output port is based upon a quantity locations in the data cell queue which are capable of receiving a data cell from an input port.

5 30. The data communication network of claim 23, wherein each output port includes logic for reassembling data frames having a destination associated with the output port from data cells received from the switching sections coupled to the output port.

SUBA 7
10 31. The data communication network of claim 30, wherein each output port is coupled to each MAC device associated with the output port through a common transmission medium and wherein each output port maintains a MAC queue of reassembled data frames for transmission to the associated MAC devices, the destination of each reassembled data frame in the MAC queue being associated with the MAC device.

15 32. The data communications network of claim 31, wherein each of the output ports transmits a signal to each of the switching sections indicating an ability to receive data cells from data links coupling the output port to the switching section.

20 33. The data communication network of claim 23, wherein for each of the output ports, each of the switching sections transmits a signal to each of the input ports indicating the ability of the switching section to receive data cells of data frames having a destination associated with the output port.